

Self-triggered Pursuit of a Single Evader with Uncertain Information [★]

Saad A. Aleem ^aCameron Nowzari ^aGeorge J. Pappas ^a^a*Department of Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA*

Abstract

This paper studies a pursuit-evasion problem involving a single pursuer and a single evader, where we are interested in developing a pursuit strategy that doesn't require continuous, or even periodic, information about the position of the evader. We propose a self-triggered control strategy that allows the pursuer to sample the evader's position autonomously, while satisfying desired performance metric of evader capture. The work in this paper builds on the previously proposed self-triggered pursuit strategy which guarantees capture of the evader in finite time with a finite number of evader samples. However, this algorithm relied on the unrealistic assumption that the evader's exact position was available to the pursuer. Instead, we extend our previous framework to develop an algorithm which allows for uncertainties in sampling the information about the evader, and derive tolerable upper-bounds on the error such that the pursuer can guarantee capture of the evader. In addition, we outline the advantages of retaining the evader's history in improving the current estimate of the true location of the evader that can be used to capture the evader with even less samples. Our approach is in sharp contrast to the existing works in literature and our results ensure capture without sacrificing any performance in terms of guaranteed time-to-capture, as compared to classic algorithms that assume continuous availability of information.

Key words: Pursuit-evasion; Self-triggered control; Sampled-data control; Set-valued analysis

1 Introduction

In this paper we study a continuous-time pursuit-evasion problem, involving a single pursuer and a single evader where the objective of the pursuer is to catch the evader. Traditionally, treatment of this problem assumes continuous or periodic availability of sensing/communication on the part of the agents, which entails numerous unwanted drawbacks like increased energy expenditure in terms of sensing requirement, network congestion, inefficient bandwidth utilization, increased risk of exposure to adversarial detection, etc. In contrast, we are interested in the scenario where we can relax this continuous/periodic sensing requirement for the pursuer and replace it with triggered decision making, where the pursuer autonomously decides when it needs to sense the evader and update its trajectory to guarantee capture of the evader.

Literature review: There are two main areas related to the contents of this paper. The first is the popular problem of pursuit-evasion which has garnered a lot of interest in the past. From an engineering perspective, pursuit-evasion problems have been studied extensively in context of differential games [Isaacs, 1999, Başar and Olsder, 1999]. In [Sgall, 2001], sufficient conditions are derived for a pursuer to capture an evader where the agents have equal maximum speeds and are constrained to move within the nonnegative quadrant of \mathbb{R}^2 . In [Alonso et al., 1992], upper and lower bounds on the time-to-capture have been discussed where the agents are constrained in a circular environment. These pursuit strategies have been generalized and extended by [Koppa and Ravishankar, 2005] to guarantee capture using multiple pursuers in an unbounded environment \mathbb{R}^n , as long as the evader is initially located inside the convex hull of the pursuers. In context of multi-agent robotic systems, the visibility-based pursuit-evasion has received a lot of interest in the past [LaValle and Hinrichsen, 2001, Sachs et al., 2004, Isler et al., 2005]. In these problems, the pursuer is visually searching for an unpredictable evader that can move arbitrarily fast in a simply connected polygonal environment. Similar problems have been studied in [Suzuki and Yamashita, 1992, Gerkey et al., 2006,

[★] The material in this paper was partially presented at the 54th IEEE Conference on Decision and Control, December 15 - 18, 2015, Osaka, Japan. Corresponding author: Saad A. Aleem. Tel.: +1-832-670-9958.

Email addresses: aleems@seas.upenn.edu (Saad A. Aleem), cnowzari@seas.upenn.edu (Cameron Nowzari), pappasg@seas.upenn.edu (George J. Pappas).

Isler et al., 2006], where visibility limitations are introduced for the pursuers but the agents actively sense and communicate at all times. A related problem has been discussed in [Bopardikar et al., 2007], where the agents can move in \mathbb{R}^2 but each agent has limited range of spatial sensing. A detailed review of recent applications in context of search and rescue missions and motion planning involving adversarial elements can be found in [Chung et al., 2011].

The vast literature on pursuit-evasion problems highlights their multifaceted applications in a variety of contexts. However, the previous works usually assume continuous or periodic availability of sensing information, especially on the part of the pursuer. Towards this end, we want to apply the new ideas of triggered control to the problem of pursuit-evasion, which has not been studied so far. In contrast to conventional time-driven approaches, strategies based on triggered control schemes study how information could be sampled for control purposes where the agents act in an opportunistic fashion to meet their desired objective [Åström and Bernhardsson, 1999, Velasco et al., 2003]. Triggered control allows us to analyze the cost to make up for less communication effort on the part of the agents, while achieving a desired task with a guaranteed level of performance of the system (see [Heemels et al., 2012] for an overview of more recent studies). Of particular relevance to this paper are works that study self-triggered [Subramanian and Fekri, 2006, Nowzari and Cortés, 2012] or event-triggered [Eqtami et al., 2010, Mazo and Tabuada, 2011] implementations of local agent strategies. In event-triggered control, the focus is on detecting events (both intrinsic and exogenous) during the execution that trigger pre-defined agent actions. In self-triggered control, the emphasis is instead on developing autonomous tests that rely only on current information available to individual agents to schedule or pre-compute future actions. In the context of pursuit-evasion problems, we will make use of the self-triggered approach which equips the pursuer with autonomous decision making in order to decrease its required sensing effort in tracking the evader. In principle, our paper shares with the above works the aim of trading increased decision making at the agent (pursuer) level for less sensing effort while still guaranteeing capture of the evader.

Contribution: This paper builds on our earlier work in [Aleem et al., 2015], where we applied the framework of triggered control to design a self-triggered pursuit policy for the pursuer which guarantees capture of the evader with a finite number of observations. Our work was different from the existing methods in the literature as our analysis did not assume the availability of continuous or periodic information about the evader. Instead, the self-triggered framework guaranteed capture of the evader without sacrificing any performance, in terms of guaranteed time-to-capture, as compared to classic algorithms that assume continuous information is available at all times.

The key result in [Aleem et al., 2015] relied on receiving perfect information about the evader, whenever the pursuer decided to sample. In reality, exact position information about the evader may never be available to the pursuer. The main contribution in this paper is to design a robust self-triggered policy, where we allow for noisy sensor measurements on the part of the pursuer, while still guaranteeing capture with only sporadic evader observations. Our triggered control framework provides fresh insights into dealing with information uncertainties and worst-case scenarios in the pursuit-evasion problem. Our framework readily incorporates the uncertainty in sensing the evader and allows us to derive tolerable error bounds in estimating the evader’s position that preserve our capture guarantees. The theme of our pursuit policy is quite similar to existing works on triggered control; based on the latest current estimate of the evader, the pursuer computes a certificate of sleep duration for which it can follow its current trajectory without having to sense the evader. In addition, we discuss the relative advantages of retaining previous estimates, where we leverage the past information about the evader to arrive at a better estimate of the evader’s true location. We show that incorporating additional knowledge about evader’s past improves the self-triggered update duration for the pursuer and mitigates uncertainty in detecting the evader.

Organization: The problem formulation and its mathematical model are presented in Section 2. In Section 3, we present the design of self-triggered update duration for the pursuer as derived in [Aleem et al., 2015]. It is followed by Section 4, where we allow for uncertainty in sampling the evader’s position and outline the maximum tolerable error that can be accommodated on the part of the pursuer, without compromising its self-triggered strategy. Section 5 discusses the relative merits of retaining previous estimates of the evader’s position in the hope of increasing the sleep durations for the pursuer. The readers are encouraged to go over the detailed analysis of our problem in the Appendix.

Notation: We let $\mathbb{R}_{>0}$, $\mathbb{R}_{\geq 0}$ and $\mathbb{Z}_{\geq 0}$ to be the sets of positive real, nonnegative real and nonnegative integer numbers, respectively. \mathbb{R}^n denotes the n -dimensional Euclidean space and $\|\cdot\|$ is the Euclidean distance.

2 Problem statement

We consider a system with a single pursuer P and a single evader E . At any given time t , the position of the evader is given by $r_e(t) \in \mathbb{R}^2$ and its velocity is given by $u_e(t) \in \mathbb{R}^2$ with $\|u_e(t)\| \leq v_e$, where v_e is the maximum speed of the evader. Similarly, the position and velocity of the pursuer are given by $r_p(t)$ and $u_p(t)$ with $\|u_p(t)\| \leq v_p$, where $v_p > v_e$ is the maximum speed

of the pursuer. The system evolves as

$$\begin{aligned}\dot{r}_p &= u_p, \\ \dot{r}_e &= u_e.\end{aligned}$$

In our problem, the goal of the pursuer is to capture the evader. We define *capture* of the evader as the instance when the pursuer is within some pre-defined capture radius $\varepsilon > 0$ of the evader. Assuming that the pursuer has exact information about the evader's state at all times, it is well known that the time-optimal strategy for the pursuer is to move with maximum speed in the direction of the evader [Isaacs, 1999]. Such a strategy, known as classical pursuit, is given by the control law

$$u_p(t) = v_p \frac{r_e(t) - r_p(t)}{\|r_e(t) - r_p(t)\|}. \quad (1)$$

The issue with the control law (1) is that it requires continuous access to the evader's state at all times and instantaneous updates of the control input. Instead, we want to guarantee capture of the evader without tracking it at all times and only updating the controller sporadically. We do this by having the pursuer decide in an opportunistic fashion when to sample the evader's position, and update its control input. Under this framework, the pursuer only knows the position of the evader at the time of its last observation. Let $\{t_k\}_{k \in \mathbb{Z}_{\geq 0}}$ be a sequence of times at which the pursuer receives information about the evader's position. In between updates, the pursuer implements a zero-order hold of the control signal computed at the last time of observation using (1) which is given by

$$u_p(t) = v_p \frac{r_e(t_k) - r_p(t_k)}{\|r_e(t_k) - r_p(t_k)\|}, \quad (2)$$

for $t \in [t_k, t_{k+1})$.

In this paper, our purpose is to identify a function for the self-triggered update duration ϕ for the pursuer that determines the next time at which the updated information is required. In other words, each time the pursuer receives updated information about the evader at some time t_k , we want to find the duration, $\phi(D_k, v_e, v_p)$, until the next update such that

$$t_{k+1} = t_k + \phi(D_k, v_e, v_p), \quad (3)$$

where $D_k \triangleq \|r_e(t_k) - r_p(t_k)\|$ is the separation between the agents at time t_k . Our goal is to design the triggering function ϕ such that the pursuer is guaranteed to capture the evader while also being aware of the number of samples of the evader required.

3 Design of Self-triggered Update Law

We study the continuous-time pursuit and evasion problem consisting of a single pursuer and a single evader

on a plane (\mathbb{R}^2), where both agents are modelled as single integrators. Note that it is not necessary to assume that the agents, particularly the evader, are moving with constant speeds at all times. For all practical purposes, we can upper bound the evader speed by v_e^{\max} such that $v_e^{\max} < v_p$ and the analysis will remain unchanged.

We denote the positions of the agents by $r_p = (x_p, y_p)$ and $r_e = (x_e, y_e)$. Additionally, the pursuer is moving along θ_p and the relative angle between the agents' headings is denoted by θ_e (see Fig. 1). Without loss of generality, we normalize the speed of the pursuer to $v_p = 1$ and the evader moves with a speed ν , where $\nu < 1$ at all times. The dynamics of the pursuer and the evader are given by

$$\begin{aligned}\dot{x}_p &= \cos \theta_p, & \dot{x}_e &= \nu \cos(\theta_e + \theta_p), \\ \dot{y}_p &= \sin \theta_p, & \dot{y}_e &= \nu \sin(\theta_e + \theta_p).\end{aligned} \quad (4)$$




Fig. 1. Figure shows the pursuer P at $r_p = (x_p, y_p)$ and the evader E at $r_e = (x_e, y_e)$ in \mathbb{R}^2 . The pursuer is moving along θ_p and the relative angle between agents' headings is denoted by θ_e . The arrows indicate the velocity vectors of the agents.

3.1 Self-triggered Update Policy for Pursuer

Suppose at time t_k the pursuer, at $r_p(t_k)$, observes the evader at $r_e(t_k)$ such that the distance between the agents is $D_k \triangleq \|r_p(t_k) - r_e(t_k)\|$. For notational brevity, we will denote the position of the agents at the instance of observation by $r_e^k := r_e(t_k)$ and $r_p^k := r_p(t_k)$. We are interested in the duration for which the pursuer can maintain its course of trajectory, without observing the evader. More specifically, we are interested in the first instance at which the separation between the agents can possibly increase, thus prompting the pursuer to sample the evader's state and update its trajectory. Let $r(t)$ denote the separation between the pursuer and the evader at time t . Then, we consider the objective function

$$R = \frac{r^2}{2} = \frac{(x_e - x_p)^2 + (y_e - y_p)^2}{2}.$$

Note that the time at which \dot{R} becomes nonnegative is same as the time at which \dot{r} becomes nonnegative. Using (4), the derivative of R (see Appendix for details) is given by

$$\dot{R} = \nu(x_e - \tau) \cos \theta_e + \nu y_e \sin \theta_e + \tau - x_e, \quad (5)$$

for $\tau \geq 0$, and is a function of evader parameters (x_e, y_e, θ_e) . For $\tau \geq 0$, the reachable set of the evader




Fig. 2. Figure shows a plot of normalized update time $\frac{\phi_\nu^k}{D_k}$ against evader speed $\nu \in [0, 1)$. $\phi_\nu^k(D_k)$ is given by (6).

is given by the ball $B_e(r_e^k, \nu\tau)$. Additionally, for fixed τ , we write \dot{R} in (5) explicitly as a function of evader parameters and denote it by $\dot{R}_\tau(x_e, y_e, \theta_e)$. Let $g(\tau) = \sup_{x_e, y_e, \theta_e} \dot{R}_\tau(x_e, y_e, \theta_e)$, subject to the reachable set of the evader, i.e. $r_e \in B_e(r_e^k, \nu\tau)$. For the dynamics in (4), where $v_p = 1$ and $v_e = \nu$, we can denote update duration in (3) by $\phi_\nu^k \triangleq \phi(D_k, v_p, v_e)$ and it is defined as

$$\phi_\nu^k = \inf \{\tau \in \mathbb{R}_{>0} \mid g(\tau) = 0\}.$$

ϕ_ν^k is the smallest duration after which there exists an evader state that may increase the separation between the agents. For the agents modelled by (4), our self-triggered update duration is obtained by solving $\inf \{\tau \in \mathbb{R}_{>0} \mid g(\tau) = 0\}$ (see Appendix for derivation) and is given by

$$\phi_\nu^k = \frac{D_k \nu \sqrt{1 - \nu^2} - D_k(1 - \nu^2)}{2\nu^2 - 1}. \quad (6)$$

The graph of ϕ_ν^k against evader speed ν is shown in Fig. 2. From the plot, we observe that increasing the evader speed decreases the self-triggered update duration for the pursuer. So, if the evader moves faster our law prescribes more frequent updates of it to guarantee capture.

The underlying objective in the design of the self-triggered policy is that at each instance of fresh observation, the separation between the agents must have decreased. The following proposition characterizes this result.

Proposition 3.1 (Decreasing Separation) *Let the pursuer and evader dynamics be given by (4) where the agents are separated by D_k at time t_k . If the pursuer updates its trajectory using the self-triggered update policy ϕ_ν^k in (6), then the distance between the agents at time t_{k+1} has strictly decreased, i.e., $D_{k+1} < D_k$ for $t_{k+1} = t_k + \phi_\nu^k$.*

PROOF. Given the separation D_k at time t_k , the new separation between the agents, after a duration of ϕ_ν^k , is

D_{k+1} . To see that the separation is strictly decreasing, note that

$$D_{k+1} \leq D_{\max}^k = D_k - (1 - \nu)\phi_\nu^k(D_k) \triangleq D_k h(\nu),$$

where D_{\max}^k is the maximum possible separation between the agents after the duration ϕ_ν^k (see Appendix for details) and $h(\nu)$ is given by

$$h(\nu) = 1 - \frac{\nu(1 - \nu)\sqrt{1 - \nu^2} - (1 - \nu)(1 - \nu^2)}{2\nu^2 - 1}, \quad (7)$$

where $h(\nu) \in [0, 1)$ for $\nu \in [0, 1)$. Thus, $D_{k+1} < D_k$. \square

3.2 Capture Time & Number of Samples

Using the self-triggered update policy (6), the pursuer is guaranteed to capture the evader in finite time with finite number updates. More specifically, we can find the maximum number of samples in terms of the capture radius ε and evader speed ν and use it to guarantee finite time-to-capture. This is summarized in the following theorem.

Theorem 3.2 (Capture with Finite Samples) *Let the pursuer and evader dynamics be given by (4) where the agents are initially separated by D_0 . Given some pre-defined positive capture radius $\varepsilon < D_0$, the self-triggered update policy ϕ_ν^k in (6) ensures capture with finite observations in finite time.*

PROOF. According to Proposition 3.1, the separation between the agents is strictly decreasing between successive updates. In fact, the new separation between the agents satisfies the inequality $D_{k+1} \leq D_k h(\nu)$, where $h(\nu)$ is given by (7). This implies that after n observations of the evader, the separation between the agents satisfies the inequality $D_n \leq D_0 h^n(\nu)$, where D_0 is the initial separation between the agents. Using this result, the maximum number of samples can be calculated by setting $D_0 h^n(\nu) \leq \varepsilon$.

$$n_{\max} = \left\lceil \frac{\log\left(\frac{\varepsilon}{D_0}\right)}{\log(h(\nu))} \right\rceil. \quad (8)$$

The expression in (8) shows that for any pre-defined positive capture radius $\varepsilon < D_0$, the pursuer is guaranteed to capture the evader with finite number of samples. This completes the first part of the proof. For self-triggered pursuit policy ϕ_ν^k in (6), the sequence of times at which the pursuer samples the evader position, denoted by $\{t_k\}_{k \in \mathbb{Z}_{\geq 0}}$, follows the criteria $t_{k+1} = t_k + \phi_\nu^k$. This means that after N updates, the total duration of pursuit (denoted by T_N) is given by $T_N = t_0 + \sum_{k=0}^N \phi_\nu^k$. Without loss of generality, we can assume $t_0 = 0$. Since the pursuer is guaranteed to capture the evader with

finite number of maximum samples n_{\max} , the time-to-capture (denoted by T_{cap}) is bounded by

$$T_{\text{cap}} \leq \sum_{k=0}^{n_{\max}} \phi_{\nu}^k = f(\nu) \sum_{k=0}^{n_{\max}} D_k,$$

where $f(\nu)$ denotes $\frac{\phi_{\nu}^k}{D_k}$ and satisfies the relationship $f(\nu) = \frac{1-h(\nu)}{1-\nu}$. Using the inequality $D_k \leq D_0 h^k(\nu)$, we get

$$T_{\text{cap}} \leq D_0 f(\nu) \sum_{k=0}^{n_{\max}} h^k(\nu).$$

As $h(\nu) \in [0, 1)$ for $\nu \in [0, 1)$, we have

$$\sum_{k=0}^{n_{\max}} h^k(\nu) < \sum_{k=0}^{\infty} h^k(\nu) = \frac{1}{1-h(\nu)}.$$

Thus, $T_{\text{cap}} < \frac{D_0}{1-\nu}$, where in the last step we have made use of the relationship $f(\nu) = \frac{1-h(\nu)}{1-\nu}$. This shows that, for any evader speed $\nu \in [0, 1)$, T_{cap} is finite. \square

Remark 3.3 (Self-triggered Performance) Note that in proving finite time-to-capture in Theorem 3.2, we showed that time-to-capture is strictly less than $\frac{D_0}{1-\nu}$. However, given a pre-defined capture radius $\varepsilon < D_0$, it can be shown that the time-to-capture satisfies the inequality

$$T_{\text{cap}} \leq \frac{D_0 - \varepsilon}{1 - \nu}. \quad (9)$$

This is because the evader is captured as soon as the actual separation is within the capture radius at any time, not just at the instance of updates. The relationship in (9) is the same upper bound for time-to-capture in classical pursuit strategy. In classical pursuit, T_{cap} is bounded by

$$T_{\text{cap}} \leq \frac{D_0 - \varepsilon}{v_p - v_e}, \quad (10)$$

where D_0 is the initial separation, ε is the pre-defined capture radius and $v_p > v_e$ [Isaacs, 1999]. The worst-case time-to-capture occurs in the scenario where evader is actively moving away at all times. This shows that our self-triggered pursuit policy guarantees capture with the *same* performance as the classical case but with only finite number of evader samples. \bullet

The expression in (8) guarantees capture with finite samples of evader's state. Fig. 3 shows a graph of the maximum number of samples required to guarantee capture against the evader speed, for pre-defined capture radius $\varepsilon = \frac{D_0}{10^3}$. The number of samples increase quite sharply as ν approaches 1. This makes intuitive sense as the maximum number of evader observations should increase as evader approaches the maximum speed of the pursuer.




Fig. 3. Figure shows the variation of the maximum number of samples against evader speed as given by the expression in (8) where ε is chosen as $\frac{D_0}{10^3}$.

4 Allowable Error in Evader's Estimate

In this section, we study the scenario in which the pursuer acquires the information about the position of the evader with some uncertainty. We are interested in analyzing the effect of imperfect observations on our self-triggered framework. Our objective is to investigate the maximum allowable error in estimating the evader's position which still allows us to catch the evader using a self-triggered pursuit policy. More specifically, we want to find the maximum tolerable uncertainty in estimating evader's position at each instance of observation, as a function of the evader's speed.

Suppose at t_k the pursuer estimates the evader at $\hat{r}_e(t_k)$. This observation is imperfect and is corrupted by an associated noise γ_k , where $\gamma_k \in \mathbb{R}_{\geq 0}$. So, at the instance of observation, the true position of the evader $r_e(t_k) \in B_e(\hat{r}_e(t_k), \gamma_k)$. For notational brevity we will denote $\hat{r}_e(t_k)$ by \hat{r}_e^k . Our objective is to find the allowable range for the error γ_k , as a function of the evader speed ν , such that the self-triggered pursuit policy can still guarantee capture. The reachable set of the evader is given by $B_e(\hat{r}_e^k, \nu(t - t_k) + \gamma_k)$, for $t \in [t_k, t_{k+1})$. This is illustrated in Fig. 4. Let $\tau = t - t_k$. Applying the previous framework of analysis, we want to maximize \dot{R} over the evader parameters, subject to the constraint $r_e \in B_e(\hat{r}_e^k, \nu\tau + \gamma_k)$. Let $g_{\gamma_k}(\tau) = \sup_{x_e, y_e, \theta_e} \dot{R}_{\tau}(x_e, y_e, \theta_e)$, subject to the reachable set of the evader, i.e. $r_e \in B_e(\hat{r}_e^k, \nu\tau + \gamma_k)$. Finding $g_{\gamma_k}(\tau)$ is very similar to the procedure of finding $g(\tau)$ as outlined in the Appendix. The only difference is in the reachable set of the evader, which is now increased by γ_k , whereas \dot{R} remains the same. For the error γ_k and estimated separation \hat{D}_k , the self-triggered update duration is defined as

$$\phi_{\gamma_k, \nu}^k = \phi(\hat{D}_k, \gamma_k, \nu) \triangleq \inf \{ \tau \in \mathbb{R}_{>0} \mid g_{\gamma_k}(\tau) = 0 \}.$$

Solving $\inf \{ \tau \in \mathbb{R}_{>0} \mid g_{\gamma_k}(\tau) = 0 \}$ yields

$$\phi_{\gamma_k, \nu}^k = \frac{\hat{D}_k \nu \sqrt{1 - \nu^2} - \hat{D}_k (1 - \nu^2) + \gamma_k (\sqrt{1 - \nu^2} - \nu)}{2\nu^2 - 1}, \quad (11)$$




Fig. 4. Figure shows the pursuer at r_p^k estimating the observer at \hat{r}_e^k , separated by \hat{D}_k , at time t_k . The pursuer detects the evader with an uncertainty γ_k . $B_e(\hat{r}_e^k, \nu(t-t_k) + \gamma_k)$ indicates the reachable set of the evader for $t \in [t_k, t_{k+1})$.

for $\nu \in [0, 1)$. To simplify the analysis, we can select the error as a scaled version of the current estimate of the separation, i.e. $\gamma_k \triangleq \beta \hat{D}_k$, where $\beta \in \mathbb{R}_{\geq 0}$. Such a parametrization will allow us to study the effect of changing β on the self-triggered update duration and will also tell us the maximum tolerable error relative to the current estimate of the separation \hat{D}_k . Setting $\gamma_k := \beta \hat{D}_k$, we get

$$\phi_{\beta, \nu}^k = \frac{\hat{D}_k (\nu \sqrt{1 - \nu^2} - (1 - \nu^2) + \beta(\sqrt{1 - \nu^2} - \nu))}{2\nu^2 - 1}, \quad (12)$$

where, with a slight abuse of notation, we have used $\phi_{\beta, \nu}^k$ instead of $\phi(\hat{D}_k, \beta, \nu) = \inf\{\tau \in \mathbb{R}_{>0} | g_\beta(\tau) = 0\}^1$. Note that β can not be chosen arbitrarily. To find the feasible domain of the of β , we invoke the criteria of $\phi_{\beta, \nu}^k \in \mathbb{R}_{>0}$ for $\nu \in [0, 1)$. This yields $\beta \in [0, \sqrt{1 - \nu^2})$. However, imposing the positivity of self-triggered duration is not sufficient to come up with the desired domain. The design of the self-triggered update duration rests on the underlying performance objective of evader capture. This requires strict decrease in true separation in between updates². At the instance of update, the new estimate of the separation satisfies $\hat{D}_{k+1} \leq D_{\max}^k$ where D_{\max}^k is given by

$$D_{\max}^k = \hat{D}_k + \nu \phi_{\beta, \nu}^k + \gamma_k + \gamma_{k+1} - \phi_{\beta, \nu}^k.$$

This is illustrated in Fig. 5. Thus,

$$\hat{D}_{k+1} \leq D_{\max}^k = \hat{D}_k + \nu \phi_{\beta, \nu}^k + \gamma_k + \gamma_{k+1} - \phi_{\beta, \nu}^k.$$

Recall that $\phi_{\beta, \nu}^k$ is a function of \hat{D}_k . Let $\bar{\phi}_{\beta, \nu}^k \triangleq \frac{\phi_{\beta, \nu}^k}{\hat{D}_k}$. This means that $\hat{D}_{k+1} \leq \hat{D}_k \left(\frac{1 - (1 - \nu) \bar{\phi}_{\beta, \nu}^k + \beta}{1 - \beta} \right)$. In deriving




Fig. 5. At time t_k , the pursuer P , measures the evader E at \hat{r}_e^k . After the duration $\phi_{\beta, \nu}^k$, \hat{D}_{k+1} indicates $\|r_p^{k+1} - \hat{r}_e^{k+1}\|$. $B_e(\hat{r}_e^k, \nu \phi_{\beta, \nu}^k + \gamma_k)$ outlines the boundary of the reachable set of the evader after $\phi_{\beta, \nu}^k$. D_{\max}^k denotes the maximum possible separation between the agents at $t_{k+1} = t_k + \phi_{\beta, \nu}^k$.




Fig. 6. Figure shows the variation of $\phi_{\beta, \nu}^k$ in (12) against β , for different values of evader speeds, where β satisfies the condition in (13). For any evader speed $\nu \in [0, 1)$, increasing the value of β decreases $\phi_{\beta, \nu}^k$.

the previous inequality, we have allowed for the worst-case scenarios in estimating the evader's position. Therefore, setting $\frac{1 - (1 - \nu) \bar{\phi}_{\beta, \nu}^k + \beta}{1 - \beta} < 1$ is sufficient to guarantee strict decrease in actual separation in between updates. This results in a more conservative set of allowable values for β as shown in (13).

$$\beta \in \left[0, \frac{(1 - \nu) \sqrt{1 - \nu^2} - 2(1 - \nu)^2}{5\nu - 3} \right). \quad (13)$$

Thus, for evader speed ν , the maximum allowable error for k -th observation (denoted by γ_k^*) satisfies the inequality

$$\gamma_k^* < \frac{\hat{D}_k ((1 - \nu) \sqrt{1 - \nu^2} - 2(1 - \nu)^2)}{5\nu - 3}.$$

Note that the maximum allowable error dynamically changes (decreases) as the pursuer closes in on the evader. This also means that, for the duration of *entire* pursuit, the maximum allowable error for *all* observa-

¹ $g_\beta(\tau) = \sup_{x_e, y_e, \theta_e} \dot{R}_\tau(x_e, y_e, \theta_e)$, s.t. $r_e \in B_e(\hat{r}_e^k, \nu\tau + \beta\hat{D}_k)$

² In the case of uncertainty, we have access to the estimates of current separation \hat{D}_k , instead of true separation D_k

tions (denoted by γ^*) satisfies the relationship

$$\gamma^* < \frac{\varepsilon((1-\nu)\sqrt{1-\nu^2}-2(1-\nu)^2)}{5\nu-3},$$

where ε is the pre-defined capture radius.

We incur no Zeno behavior using the self-triggered policy in (12), i.e. the pursuer does not require infinitely many samples to capture the evader. The following theorem characterizes this important result.

Theorem 4.1 (No Zeno Behavior) *If the pursuer updates its trajectory using the self-triggered update policy $\phi_{\beta,\nu}^k$ in (12), where β satisfies the condition in (13), then pursuer is guaranteed to incur no Zeno behavior for the duration of the pursuit.*

PROOF. Suppose that pursuer observes the evader at the instance t_k . In order to show no Zeno behavior, it suffices to prove that the inter-event duration is lower-bounded by a positive constant for all observations, i.e. $t_{k+1} - t_k := \phi_{\beta,\nu}^k \geq c > 0, \forall k$. According the condition (13), $\beta \in \left[0, \frac{(1-\nu)\sqrt{1-\nu^2}-2(1-\nu)^2}{5\nu-3}\right)$ for the self-triggered update policy $\phi_{\beta,\nu}^k$ in (12). Note that the higher values of β result in a smaller inter-event duration (see Fig. 6). Then, given a capture radius ε , the following relationship is satisfied for all observations.

$$\phi_{\beta,\nu}^k \geq \phi(\varepsilon, \nu) = \varepsilon q(\nu),$$

where $q(\nu) = \frac{\nu\sqrt{1-\nu^2}-(1-\nu^2)+p(\nu)(\sqrt{1-\nu^2}-\nu)}{2\nu^2-1}$ and $p(\nu) = \frac{(1-\nu)\sqrt{1-\nu^2}-2(1-\nu)^2}{5\nu-3}$. It is easy to verify that $q(\nu)$ is a positive and monotonically decreasing function for $\nu \in [0, 1)$. Thus, for any evader speed $\nu \in [0, 1)$, we have $t_{k+1} - t_k := \phi_{\beta,\nu}^k \geq \varepsilon q(\nu) > 0, \forall k$. So, inter-event duration is lower-bounded by a positive constant for all observations, which suffices to show that our self-triggered duration does not incur any Zeno behavior. \square

An important consequence of Theorem 4.1 is that, in the presence of uncertainty, our self-triggered framework guarantees capture with only finite estimates of the evader. This means that we can find the maximum number of samples that will guarantee capture. By design, β satisfies the condition in (13) and guarantees strict decrease in measured separation between successive updates. In general, we can write this as

$$\hat{D}_{k+1} \leq \hat{D}_k h_\beta(\nu),$$

where $h_\beta(\nu)^1 \in [0, 1)$ for $\nu \in [0, 1)$ and β satisfies the condition in (13). Using similar analysis from previous section, we can estimate the maximum number of

¹ $h_\beta(\nu) = \frac{1 - \frac{(1-\nu)(\nu\sqrt{1-\nu^2}-(1-\nu^2)+\beta(\sqrt{1-\nu^2}-\nu))}{2\nu^2-1}}{1-\beta}$.

samples in the presence of uncertainty. Given an initial estimated separation of \hat{D}_0 between agents and a pre-defined positive capture radius $\varepsilon < \hat{D}_0$, the maximum number of evader updates can be calculated by setting $\hat{D}_0 (h_\beta(\nu))^n \leq \varepsilon$, which results in

$$n_{\max}^\beta = \left\lceil \frac{\log\left(\frac{\varepsilon}{\hat{D}_0}\right)}{\log(h_\beta(\nu))} \right\rceil. \quad (14)$$

For evader speed ν , the maximum number of samples for guaranteed capture *increases* as the parameter β is increased (see Fig. 7). This shows that, for any $\nu \in [0, 1)$, we need to sample the evader's state more frequently in order to allow for more uncertainty in estimating evader's position at each instance of update.




Fig. 7. Figure shows the plot of maximum number of samples in (14) against β , for different values of evader speeds, where β satisfies the condition in (13). The capture radius is taken as $\varepsilon = \frac{\hat{D}_0}{10^3}$.

5 Self-triggered Analysis with Memory

In the previous section, we studied the effect of uncertainty in sensing the evader's position and provided bounds on the maximum tolerable error (as a function of evader's speed) which allowed us to capture the evader with sporadic updates. In the absence of uncertainty, as outlined in Section 3, the pursuer employs a *memoryless* pursuit policy to catch the evader, i.e. the pursuer only needs the *current* sample of the evader's (true) position r_e^k in order to calculate the self-triggered update duration ϕ_ν^k in (6). This is because, for $\tau \in [0, \phi_\nu^k)$, the *current* reachable set of the evader, $B_e(r_e^k, \nu\tau)$, is always a subset of the *previous* reachable set of the evader, $B_e(r_e^{k-1}, \nu(\phi_\nu^k + \tau))$. When we introduce uncertainty in estimating evader's position, the above statement is no longer true. The idea behind retaining evader's estimated history is that we can potentially reduce the *actual* reachable set of the evader when we combine the current reachable set of the evader with the previous reachable sets, thus improving our current estimate of the true position of the evader. This allows us to improve (increase) our update duration, as compared to

the memoryless case *with* uncertainty (derived in Section 4). More specifically, we can leverage our knowledge about the previous estimates of the evader in improving the current update duration for the pursuer, while mitigating the effect of uncertainty in estimating evader's position.

Consider the problem in which the pursuer receives uncertain information about the evader, while keeping track of its previous estimates. For the purpose of illustration, we analyze the case where the pursuer retains only the previous estimate of the evader's position. Extending the framework to the case for more than one previous estimates will be similar and straightforward. Additionally, we assume that the pursuer samples the evader with an associated error γ . Suppose that the pursuer sampled the evader at time t_{k-1} , computed the self-triggered update duration $\hat{\phi}_{\gamma,\nu}^{k-1}$ and observed the evader again at the instance t_k . For $t \in [t_k, t_{k+1})$, the *current* reachable set of the evader is given by $B_e(\hat{r}_e^k, \nu(t - t_k) + \gamma)$, and the *previous* reachable set of the evader is given by $B_e(\hat{r}_e^{k-1}, \nu(t - t_k + \hat{\phi}_{\gamma,\nu}^{k-1}) + \gamma)$. Based on this information, the *actual* reachable set of the evader is given by the intersection of the two. One particular scenario is illustrated in Fig. 8, where the actual reachable set is not the same as the current reachable set.




Fig. 8. The figure shows the reachable sets of the evader, based on its current (\hat{r}_e^k) and previous (\hat{r}_e^{k-1}) estimates. For $t \in [t_k, t_{k+1})$, the current reachable set is denoted by $B_e(\hat{r}_e^k, \nu(t - t_k) + \gamma)$ and the previous reachable set is denoted by $B_e(\hat{r}_e^{k-1}, \nu(t - t_k + \hat{\phi}_{\gamma,\nu}^{k-1}) + \gamma)$. The actual reachable set is the intersection of the two.

Using $\tau = t - t_k$ for $t \in [t_k, t_{k+1})$, let $\hat{g}_\gamma(\tau) = \sup_{x_e, y_e, \theta_e} \dot{R}_\tau(x_e, y_e, \theta_e)$, subject to the *actual* reachable set of the evader i.e.

$$r_e \in B_e(\hat{r}_e^k, \nu\tau + \gamma) \cap B_e(\hat{r}_e^{k-1}, \nu(\tau + \hat{\phi}_{\gamma,\nu}^{k-1}) + \gamma).$$

The self-triggered update duration is then defined as

$$\hat{\phi}_{\gamma,\nu}^k = \inf \{ \tau \in \mathbb{R}_{>0} \mid \hat{g}_\gamma(\tau) = 0 \}. \quad (15)$$

Equivalently, for fixed τ and γ , $\hat{g}_\gamma(\tau)$ is the optimal value of the following optimization problem.

$$\begin{aligned} & \sup_{x_e, y_e, \theta_e} \dot{R}_\tau(x_e, y_e, \theta_e), \\ & \text{subject to } r_e \in B_e(\hat{r}_e^{k-1}, \nu(\tau + \hat{\phi}_{\gamma,\nu}^{k-1}) + \gamma), \\ & r_e \in B_e(\hat{r}_e^k, \nu\tau + \gamma). \end{aligned} \quad (16)$$

As explained earlier, incorporating evader's previous estimate can potentially reduce the *actual* reachable set of the evader in the case of noisy measurements. By keeping track of the previous estimate(s), we are equivalently adding more constraints to the feasible (reachable) set in our optimization problem. We want to formalize the benefit of retaining evader's history in terms of improvement (increase) in the update duration, as compared to the memoryless update duration given by (11). Recall that the self-triggered update duration in (11) was derived using only the current reachable set of the evader, $B_e(\hat{r}_e^k, \nu\tau + \gamma)$, based on its latest estimate. Let $g_\gamma(\tau)$ denote the optimal value of the problem (17), which is used to obtain the memoryless update duration in (11).

$$\begin{aligned} & \sup_{x_e, y_e, \theta_e} \dot{R}_\tau(x_e, y_e, \theta_e), \\ & r_e \in B_e(\hat{r}_e^k, \nu\tau + \gamma). \end{aligned} \quad (17)$$

Observe that (17) is a relaxation of (16), as it is obtained by removing the constraint corresponding to the previous reachable set of the evader. As a result, $\hat{g}_\gamma(\tau) \leq g_\gamma(\tau)$ for any $\tau \geq 0$. We notice that both \hat{g}_γ and g_γ are monotonically increasing in the parameter τ , because increasing τ increases the feasible set in the optimization problem thus yielding a potentially greater maximum value. Using the monotonicity of \hat{g} and g in τ , along with the fact that $\hat{g}_\gamma(\tau) \leq g_\gamma(\tau)$ (for any $\tau \geq 0$), allows us to infer that the first instance, at which \hat{g} approaches 0, will be greater than or equal to the first instance at which g approaches 0. This means that

$$\inf \{ \tau \in \mathbb{R}_{>0} \mid \hat{g}_\gamma(\tau) = 0 \} \geq \inf \{ \tau \in \mathbb{R}_{>0} \mid g_\gamma(\tau) = 0 \}$$

and as a consequence, we have

$$\hat{\phi}_{\gamma,\nu}^k \geq \phi_{\gamma,\nu}^k. \quad (18)$$

Leveraging Memory against Uncertainty: The relationship in (18) shows that by using previous estimate(s) of the evader's position, we can potentially increase the update durations for the pursuer. Case I in Fig. 9 demonstrates one particular instance of leveraging evader's history which yields the greatest improvement




Fig. 9. Figure shows the two possible cases at the instance of fresh observation of the evader. Case I shows the extreme case where incorporating the evader's history precisely determines its true position. Case II shows the scenario in which the current reachable set is a subset of the previous reachable set, thus the previous estimate provides no additional information.

in the current update duration. In the extreme case, at the instance of observation, the current and the previous reachable sets of the evader intersect such that the *actual* reachable set is reduced to a point. Equivalently, this means that we know precisely where the evader is. This will result in a longer update duration as compared to memoryless case (11), where we would have incorporated uncertainty in our observation to yield a more conservative update duration. Let $\Delta\phi$ denote the improvement in self-triggered update duration, i.e. $\Delta\phi := |\hat{\phi}_{\gamma,\nu}^k - \phi_{\gamma,\nu}^k|$. Comparing Eq. (6) and Eq. (11), we see that the greatest improvement (denoted by $\Delta\phi^*$) is given by

$$\Delta\phi^* = \frac{2\gamma(\nu - \sqrt{1 - \nu^2})}{2\nu^2 - 1}, \quad (19)$$

for $\nu \in [0, 1)$.

While sometimes adding information about the previous estimate can be advantageous, it is important to realize that incorporating the previous reachable set of the evader does not always results in an improvement (increase) in the update duration. We can increase the self-triggered update duration only when the evader's past provides more information about its current true position. In the scenario where the current reachable set of the evader is a subset of the previous reachable set, we get no additional information about the evader's true position and hence no improvement in our update duration as compared to the memoryless case in (11). This is illustrated in case II in Fig. 9, where we can drop (forget) the previous estimate of the evader's position as it is a subset of the previous reachable set and will yield no improvement in increasing the current update duration.

Remark 5.1 (Numerical Example) To illustrate the improvement in update duration numerically, let $\nu = 0.5$ and $\gamma = 0.1$. For these values, $\Delta\phi^* = 0.146$ from (19). Based on some initial measurement, the pursuer finds the first update duration $\phi_{\gamma,\nu}^0$ by using (11) and samples the evader at t_1 to find that the previ-

ous and the current reachable set intersect at a point. Suppose the measured separation at t_1 is $\hat{D}_1 = 10$ units. Using a memoryless pursuit policy in (11) results in $\phi_{\gamma,\nu}^1 = 6.529$, whereas obtaining self-triggered update duration from solving the problem (16) results in $\hat{\phi}_{\gamma,\nu}^1 = 6.674$. Observe that $\hat{\phi}_{\gamma,\nu}^1 - \phi_{\gamma,\nu}^1 \approx \Delta\phi^*$, which shows that, in certain cases, knowing one previous estimate of the evader can almost nullify the uncertainty in sensing evader's position and consequently allow for greater update duration. •

Towards Multiple Estimates: Extending the above framework to the case of multiple previous estimates is relatively straightforward. Suppose, for the k -th observation of the evader, we have the information about the m previous updates (where $m \in \mathbb{Z}_{\geq 0}$ such that $m \leq k-1$). Then the current sleep duration can be computed by $\hat{\phi}_{\gamma,\nu}^k = \inf \{\tau \in \mathbb{R}_{>0} \mid \hat{g}_\gamma(\tau) = 0\}$, where $\hat{g}_\gamma(\tau)$ is an optimal value of the optimization problem (20) for $\tau \geq 0$. The value m can be treated as the length of the sliding window for retaining fixed number of previous estimates, to compute the current update duration for the pursuer.

$$\begin{aligned} & \sup_{x_e, y_e, \theta_e} \dot{R}_\tau(x_e, y_e, \theta_e), \\ & \text{subject to } r_e \in B_e(\hat{r}_e^k, \nu\tau + \gamma), \\ & r_e \in \bigcap_{j=1}^m B_e^j(\tau), \end{aligned} \quad (20)$$

where $B_e^j(\tau)$ is the current reachable set of the j -th previous estimate of the evader's position and is given by

$$B_e^j(\tau) := B_e\left(\hat{r}_e^{k-j}, \nu\left(\tau + \sum_{l=1}^j \hat{\phi}_{\gamma,\nu}^{k-l}\right) + \gamma\right),$$

for $j \in \{1, \dots, m\}$ and $m \leq k-1$.

Remark 5.2 (Forgetting Previous Estimates)

Note that, while we might have the capability to store m previous estimates of the evader, it is not necessary to use all of them in computing the current sleep duration for the pursuer. As illustrated earlier, retaining the history improves our update duration when it reduces the current reachable set of the evader. This allows us to forget all those estimates whose reachable sets either completely contain the current reachable set or a reachable set of another previous estimate. To formalize this notion, at the instance of k -th observation of the evader, we can construct a set

$$\mathcal{I} := \{i \mid B_e(\hat{r}_e^k, \gamma) \subseteq B_e^i(0)\} \cup \{i \mid B_e^l(0) \subseteq B_e^i(0), l \neq i\},$$

where $B_e^i(0)$ denotes $B_e\left(\hat{r}_e^{k-i}, \nu \sum_{l=1}^i \hat{\phi}_{\gamma,\nu}^{k-l} + \gamma\right)$ for $i \in \{1, \dots, m\}$ and $m \leq k-1$. \mathcal{I} denotes the collection of

indices among the m previous samples of all those estimates which we can forget to reduce the computation complexity of the problem (20). Thus, our improved update duration can be computed from solving the problem (21). The problem (21) is guaranteed to have the same optimal value as that of (20) because all the estimates belonging to the set \mathcal{I} have no effective contribution to the actual reachable set of the evader and thus removing them will have no change in the optimal value.

$$\begin{aligned} & \sup_{x_e, y_e, \theta_e} \dot{R}_\tau(x_e, y_e, \theta_e), \\ & \text{subject to } r_e \in B_e(\hat{r}_e^k, \nu\tau + \gamma), \\ & \quad r_e \in \bigcap_{j=1, j \notin \mathcal{I}}^m B_e^j(\tau). \end{aligned} \quad (21)$$

6 Simulations

In this section, we provide numerical results for the case when the pursuer retains past samples of the evader as outlined in Section 5. We study the potential benefit of retaining only the previous estimate of the evader's observation as the pursuer tries to capture the evader. Thus, in our simulations, $m = 1$ in (21) and the self-triggered update duration will be obtained from solving the optimization problem in (16).

We model our agents as single integrators, where we normalize the speeds of the agents, such that $v_p = 1$ and the maximum speed of the evader is $\nu = 0.5$. The evader is restricted to move in any of 4 directions: right, left, up and down and chooses the best direction to actively move away, with its maximum possible speed, from the pursuer at all times. We initialize the agents with an actual separation $D_0 = 15$ units. For every observation, the pursuer samples the evader's current position with an associated error of $\gamma = 0.1$ such that, initially, the true position of the evader $r_e^k \in B_e(\hat{r}_e^k, 0.1)$. Note that we can not arbitrarily set the capture radius, as outlined in Section 4. For an evader speed ν and error γ , the capture radius must satisfy the relationship $\varepsilon > \frac{\gamma}{p(\nu)}$, where

$p(\nu) = \frac{(1-\nu)\sqrt{1-\nu^2}-2(1-\nu)^2}{5\nu-3}$. For $\nu = 0.5$ and $\gamma = 0.1$, setting $\varepsilon = 0.75$ satisfies the aforementioned relationship. Table 1 shows the variation between the *memoryless* update duration $\phi_{\gamma,\nu}^k$ and the improved update duration $\hat{\phi}_{\gamma,\nu}^k$, which takes into account the previous estimate of the evader's observation. Since, $\hat{\phi}_{\gamma,\nu}^k$ is different from $\phi_{\gamma,\nu}^k$, they result in different measures of separation at the instance of observation. For a fair comparison, we need to compare the normalized update durations, i.e. $\frac{\phi_{\gamma,\nu}^k}{D_k}$ and $\frac{\hat{\phi}_{\gamma,\nu}^k}{\bar{D}_k}$ where \bar{D}_k and D_k denote the true separations at the instance of evader observation for memory-aware and memoryless pursuit strategies, respectively. The results in Table 1 indicate that in the presence of uncertainty in sampling the evader, incorporating only the previous estimate allows for greater sleep durations

k	D_k	$\phi_{\gamma,\nu}^k$	$\frac{\phi_{\gamma,\nu}^k}{D_k}$	\bar{D}_k	$\hat{\phi}_{\gamma,\nu}^k$	$\frac{\hat{\phi}_{\gamma,\nu}^k}{\bar{D}_k}$
0	15.00	9.500	0.633	15.00	9.500	0.633
1	10.25	6.488	0.633	10.24	6.674	0.652
2	7.006	4.432	0.633	6.949	4.511	0.649
3	4.790	3.027	0.632	4.659	2.996	0.643
4	3.277	2.067	0.631	3.196	2.037	0.637
5	2.243	1.412	0.630	2.194	1.393	0.634
6	1.537	0.965	0.628	1.500	0.947	0.632
7	1.055	0.659	0.625	1.033	0.648	0.628
8	0.754	0.463	0.621	0.752	0.464	0.623

Table 1
Comparison between memoryless and memory-aware update times.

that guarantees capture with finite updates. We observe that, in the beginning of the pursuit, adding history results in relatively better gains as compared to towards the end, when the agents are nearby. The values indicate that as the separation decreases, so does the potential benefit of adding computational overhead by retaining the previous observation.

7 Conclusions

The robust self-triggered framework in this paper extends our previous results to address the practical issues related with uncertainty in information about the evader. We elaborate the case when the sampling is not perfect and design self-triggered update duration along with tolerable error bounds in estimating the evader's state. We show that our analysis preserve the self-triggered controller updates for the pursuer, such that it incurs no Zeno behavior in catching the evader without losing any of the previous performance guarantees. Our methodology offers a fresh perspective on dealing with uncertain information in pursuit-evasion problems, besides being in contrast to a majority of previous works that assume continuous, or at least periodic, information about the evader is available at all times. Additionally, we study the merits of retaining evader's history and show that we can allow for potentially longer update durations by incorporating past observations of the evader in the pursuer's autonomous decision making. In the future, we are interested in extending our methods to scenarios involving multiple agents and deriving conditions for team-triggered cooperative strategies.

References

- S. A. Aleem, C. Nowzari, and G. J. Pappas. Self-triggered pursuit of a single evader. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1433–1440, Osaka, Japan, 2015.
- L. Alonso, A. S. Goldstein, and E. M. Reingold. Lion and man: upper and lower bounds. *ORSA Journal on Computing*, 4(4): 447–452, 1992.
- K. J. Åström and B. Bernhardsson. Comparison of periodic and event based sampling for first-order stochastic systems. In *Proceedings of the 14th IFAC World congress*, volume 11, pages 301–306, 1999.

- T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory*, volume 23. SIAM, 1999.
- S. D. Bopardikar, F. Bullo, and J. P. Hespanha. Cooperative pursuit with sensing limitations. In *Proceedings of the IEEE American Control Conference*, pages 5394–5399, New York, NY, 2007.
- T. H. Chung, G. A. Hollinger, and V. Isler. Search and pursuit-evasion in mobile robotics - a survey. *Autonomous Robots*, 31(4):299–316, 2011.
- A. Eqtami, D. V. Dimarogonas, and K. J. Kyriakopoulos. Event-triggered control for discrete-time systems. In *Proceedings of the IEEE American Control Conference*, pages 4719–4724, Baltimore, MD, 2010.
- B. P. Gerkey, S. Thrun, and G. Gordon. Visibility-based pursuit-evasion with limited field of view. *The International Journal of Robotics Research*, 25(4):299–315, 2006.
- W. P. M. H. Heemels, K. H. Johansson, and P. Tabuada. An introduction to event-triggered and self-triggered control. In *Proceedings of the IEEE Conference on Decision and Control*, pages 3270–3285, Maui, HI, 2012.
- R. Isaacs. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. Courier Corporation, 1999.
- V. Isler, S. Kannan, and S. Khanna. Randomized pursuit-evasion in a polygonal environment. *IEEE Transactions on Robotics*, 21(5):875–884, 2005.
- V. Isler, S. Kannan, and S. Khanna. Randomized pursuit-evasion with local visibility. *SIAM Journal on Discrete Mathematics*, 20(1):26–41, 2006.
- S. Kopparty and C. V. Ravishankar. A framework for pursuit evasion games in \mathbb{R}^n . *Information Processing Letters*, 96(3):114–122, 2005.
- S. M. LaValle and J. E. Hinrichsen. Visibility-based pursuit-evasion: The case of curved environments. *IEEE Transactions on Robotics and Automation*, 17(2):196–202, 2001.
- M. Mazo and P. Tabuada. Decentralized event-triggered control over wireless sensor/actuator networks. *IEEE Transactions on Automatic Control*, 56(10):2456–2461, 2011.
- C. Nowzari and J. Cortés. Self-triggered coordination of robotic networks for optimal deployment. *Automatica*, 48(6):1077–1087, 2012.
- S. Sachs, S. M. LaValle, and S. Rajko. Visibility-based pursuit-evasion in an unknown planar environment. *The International Journal of Robotics Research*, 23(1):3–26, 2004.
- J. Sgall. Solution of David Gale’s lion and man problem. *Theoretical Computer Science*, 259(1-2):663–670, 2001.
- R. Subramanian and F. Fekri. Sleep scheduling and lifetime maximization in sensor networks: fundamental limits and optimal solutions. In *Symposium on Information Processing in Sensor Networks*, pages 218–225, New York, NY, 2006. ACM.
- I. Suzuki and M. Yamashita. Searching for a mobile intruder in a polygonal region. *SIAM Journal on Computing*, 21(5):863–888, 1992.
- M. Velasco, J. Fuertes, and P. Marti. The self triggered task model for real-time control systems. In *Proceedings of the 24th IEEE Real-Time Systems Symposium*, volume 384, pages 67–70, 2003.

A Derivation of Self-triggered Update Duration

The agents are modelled by the dynamics in (4). At time t_k , the pursuer observes the evader at a distance $D_k = \|r_e^k - r_p^k\|$. Without loss of generality, we make the relative vector between pursuer and the evader parallel to the x -axis such that $y_p(t_k) = 0$ and $y_e(t_k) = 0$. Additionally, as a matter of convenience, we assume that $0 = x_p(t_k) < x_e(t_k) = D_k$. This is elaborated in Fig. A.1.




Fig. A.1. Figure shows the pursuer and the evader, at $r_p^k = (0, 0)$ and $r_e^k = (D_k, 0)$ respectively, separated by D_k at time t_k . $B_e(r_e^k, \nu(t - t_k))$ is the ball centered at r_e^k with radius $\nu(t - t_k)$ and indicates the reachable set of the evader for $t \in [t_k, t_{k+1})$.

Thus, our pursuit trajectory is parallel to the x -axis as $\theta_p(t_k) = 0$. The pursuer does not observe the evader till the next update instance, thus $\theta_p(t) = 0$ for $t_{k+1} - t_k$. The modified dynamics are given by

$$\begin{aligned} \dot{x}_p &= 1, & \dot{x}_e &= \nu \cos \theta_e, \\ \dot{y}_p &= 0, & \dot{y}_e &= \nu \sin \theta_e. \end{aligned} \quad (\text{A.1})$$

From (A.1), $r_p(t) = (t - t_k, 0)$ for $t \in [t_k, t_{k+1})$. Thus, $\dot{R} = (x_e - x_p)(\dot{x}_e - \dot{x}_p) + (y_e - y_p)(\dot{y}_e - \dot{y}_p)$, where $R = \frac{r^2}{2}$ and r is the separation. $\dot{R} = \nu(x_e - t + t_k) \cos \theta_e + \nu y_e \sin \theta_e + t - t_k - x_e$. Using $\tau \triangleq t - t_k$,

$$\dot{R} = \nu(x_e - \tau) \cos \theta_e + \nu y_e \sin \theta_e + \tau - x_e. \quad (\text{A.2})$$

The agent updates when \dot{R} can possibly become nonnegative. For $\tau \geq 0$, $r_e \in B_e(r_e^k, \nu\tau)$. $\dot{R}_\tau(x_e, y_e, \theta_e)$ explicitly denotes \dot{R} in (A.2) in terms of evader parameters, for fixed τ . The problem is formulated as

$$\begin{aligned} \sup_{x_e, y_e, \theta_e} & \quad \dot{R}_\tau(x_e, y_e, \theta_e), \\ \text{subject to} & \quad (x_e, y_e) \in B_e(r_e^k, \nu\tau). \end{aligned} \quad (\text{A.3})$$

For $\tau \geq 0$, $g(\tau)$ denotes the optimal value of problem (A.3) and the self-triggered update duration is defined as

$$\phi_\nu^k = \inf \{ \tau \in \mathbb{R}_{>0} \mid g(\tau) = 0 \}.$$

In \mathbb{R}^2 , $B_e(r_e^k, \nu\tau) \triangleq (x_e - D_k)^2 + y_e^2 \leq (\nu\tau)^2$. The constraint of the problem (A.3) is independent of θ_e . $\frac{\partial \dot{R}_\tau}{\partial \theta_e} = -\nu(x_e - \tau) \sin \theta_e + \nu y_e \cos \theta_e$. Setting $\frac{\partial \dot{R}_\tau}{\partial \theta_e} = 0$ yields $\theta_e^* = \arctan\left(\frac{y_e}{x_e - \tau}\right)$. Note that for $\tau \in [0, \phi_\nu^k)$, $x_e - \tau \geq 0$. To see this, suppose $x_e - \tau < 0$. Setting $\theta_e = 0$ in (A.2), we get $\dot{R}_\tau = -(1 - \nu)(x_e - \tau) > 0$ for $\nu \in [0, 1)$, which is a contradiction as $\dot{R} \leq 0$ for any $\tau \in [0, \phi_\nu^k)$. Due to symmetry of the problem, we can assume $y_e \geq 0$. Since $x_e - \tau \geq 0$, we have $\theta_e^* \in [0, \frac{\pi}{2}]$. Substituting θ_e^* in (A.2), we get

$$\dot{R}_\tau(x_e, y_e) = \nu \sqrt{y_e^2 + (x_e - \tau)^2} + \tau - x_e,$$

which simplifies the problem in (A.3) to

$$\sup_{x_e, y_e} \dot{R}_\tau(x_e, y_e) = \nu \sqrt{y_e^2 + (x_e - \tau)^2} + \tau - x_e \quad (\text{A.4})$$

subject to $(x_e, y_e) \in B_e(r_e^k, \nu\tau)$.

From the constraint $(x_e, y_e) \in B_e(r_e^k, \nu\tau)$ we get $y_e^2 \leq (\nu\tau)^2 - (x_e - D_k)^2$, which means that y_e^* must lie at the boundary of $B_e(r_e^k, \nu\tau)$. Thus, $y_e^* = \sqrt{(\nu\tau)^2 - (x_e - D_k)^2}$. Substituting y_e^* in $\dot{R}_\tau(x_e, y_e)$, we get

$$\dot{R}_\tau(x_e) = \nu \sqrt{(x_e - \tau)^2 + (\nu\tau)^2 - (x_e - D_k)^2} + \tau - x_e.$$

This reduces the problem in (A.4) to

$$\begin{aligned} \sup_{x_e} \quad & \dot{R}_\tau(x_e) \\ \text{subject to} \quad & x_e \in [D_k - \nu\tau, D_k + \nu\tau]. \end{aligned} \quad (\text{A.5})$$

Note that we can relax the problem in (A.5) by omitting the constraint $x_e \in [D_k - \nu\tau, D_k + \nu\tau]$. The relaxation of (A.5) results in an unconstrained optimization problem. Ignoring the constraints of problem (A.5), let $\tilde{g}(\tau) \triangleq \sup_{x_e} \dot{R}_\tau(x_e)$ and $\tilde{\phi}_\nu^k = \inf\{\tau \in \mathbb{R}_{>0} \mid \tilde{g}(\tau) = 0\}$.

As a result of this relaxation, we have $\tilde{\phi}_\nu^k \leq \phi_\nu^k$. Let $\tilde{x}_e^* := \operatorname{argmax} \dot{R}_\tau(x_e)$ for the unconstrained problem. To perform unconstrained maximization of $\dot{R}_\tau(x_e)$, the derivative is given by $\frac{\partial \dot{R}_\tau}{\partial x_e} = \frac{\nu(D_k - \tau)}{\sqrt{(x_e - \tau)^2 + (\nu\tau)^2 - (x_e - D_k)^2}} - 1$. Setting $\frac{\partial \dot{R}_\tau}{\partial x_e} = 0$ yields $\tilde{x}_e^* = \frac{D_k^2 \nu^2 + D_k^2 - 2\tau D_k \nu^2 - \tau^2}{2(D_k - \tau)}$. As $\tilde{x}_e^* := \operatorname{argmax} \dot{R}_\tau(x_e)$, $\tilde{g}(\tau) = \dot{R}_\tau(\tilde{x}_e^*)$ and is given by $\tilde{g}(\tau) = \tau - \frac{D_k^2 \nu^2 + D_k^2 - 2\tau D_k \nu^2 - \tau^2}{2(D_k - \tau)} + \nu^2(D_k - \tau)$. Solving for $\inf\{\tau \in \mathbb{R}_{>0} \mid \tilde{g}(\tau) = 0\}$, yields

$$\tilde{\phi}_\nu^k = \frac{D_k \nu \sqrt{1 - \nu^2} - D_k(1 - \nu^2)}{2\nu^2 - 1}. \quad (\text{A.6})$$

Recall that $\tilde{\phi}_\nu^k \leq \phi_\nu^k$, where $\tilde{\phi}_\nu^k$ was obtained from relaxing the constraint in problem (A.5). Our claim is that at the instance of update ($\tau = \tilde{\phi}_\nu^k$), the maximizer \tilde{x}_e^* of the relaxed problem is a feasible solution of the problem (A.5), i.e. $\tilde{x}_e^*(\tau) \in [D_k - \nu\tau, D_k + \nu\tau]$ for $\tau = \tilde{\phi}_\nu^k$. To see this, at $\tau = \tilde{\phi}_\nu^k$, the maximizer \tilde{x}_e^* , $D_k + \nu\tau$ and $D_k - \nu\tau$ are given by

$$\begin{aligned} D_k - \nu\tilde{\phi}_\nu^k &= D_k \varphi_-(\nu), \\ \tilde{x}_e^*(\tilde{\phi}_\nu^k) &= D_k \varphi_2(\nu), \\ D_k + \nu\tilde{\phi}_\nu^k &= D_k \varphi_+(\nu), \end{aligned}$$

where

$$\begin{aligned} \varphi_\pm(\nu) &= 1 \pm \frac{\nu^2 \sqrt{1 - \nu^2} - \nu(1 - \nu^2)}{2\nu^2 - 1}, \\ \varphi_2(\nu) &= \frac{3\nu^3 + \sqrt{1 - \nu^2} - 2\nu - 2\nu^4 \sqrt{1 - \nu^2}}{(\nu - \sqrt{1 - \nu^2})(2\nu^2 - 1)}. \end{aligned}$$

For $\nu \in [0, 1]$, $\varphi_-(\nu) \leq \varphi_2(\nu) \leq \varphi_+(\nu)$. This shows that $\tilde{x}_e^*(\tilde{\phi}_\nu^k)$ satisfies the constraints in the problem (A.5) at $\tau = \tilde{\phi}_\nu^k$. This means, at the instance of update, the maximizer \tilde{x}_e^* of the relaxed problem is a feasible solution of the original problem (A.5) and hence it is optimal solution x_e^* for (A.5). Thus $\tilde{g}(\tau) = g(\tau)$ and as a result $\tilde{\phi}_\nu^k = \phi_\nu^k$. Note that ϕ_ν^k is continuous in the parameter ν as $\lim_{\nu \rightarrow \frac{1}{\sqrt{2}}} \frac{D_k \nu \sqrt{1 - \nu^2} - D_k(1 - \nu^2)}{2\nu^2 - 1} = \frac{D_k}{2}$.

B Maximum Separation

If the pursuer updates its trajectory using the self-triggered update policy described in (A.6), then the maximum distance between the agents, between successive updates, is given by

$$D_{\max}^k = D_k - (1 - \nu)\phi_\nu^k.$$

To see this, after the duration ϕ_ν^k , the pursuer moves a distance of ϕ_ν^k units and the evader can be anywhere inside a ball of radius $\nu\phi_\nu^k$ centered at r_e^k . This is shown in Fig. B.1. The maximum separation between the pursuer and the evader is denoted by D_{\max}^k and is given by $D_k - (1 - \nu)\phi_\nu^k$.




Fig. B.1. D_{k+1} indicates the new separation between the agents at t_{k+1} . $B_e(r_e^k, \nu\phi_\nu^k)$ outlines the boundary of the reachable set of the evader after ϕ_ν^k . D_{\max}^k denotes the maximum possible separation between the pursuer and the evader at $t = t_k + \phi_\nu^k$.